July - August 1995
Volume 2, Number 11

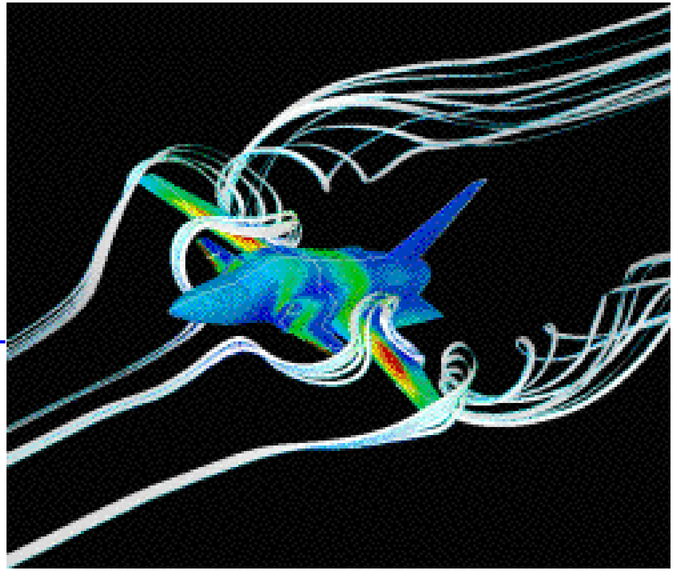

Large, Out-of-Core Calculation Runs on the IBM SP2

By Charbel Farhat

Philippe Homsi and Nicolas Durante contributed to this article.

Researchers at the University of Colorado at Boulder recently developed a portable out-of-core direct dense solver that tackles electromagnetic and elastic scattering problems. Using this solver and the IBM SP2 parallel system at NAS, a team in Boulder, working closely with NAS, ran a substantial out-of-core calculation. The project demonstrated the feasibility of simulating electromagnetic scattering on a moderately sized, target benchmark problem.

The systems of equations associated with the boundary element formulations of these problems involve generally dense complex matrices whose size is dictated by the specified wave number to the fourth power. As a result, scattering simulations in the gigahertz (GHz) regime -- which typically generate systems with more than 100,000 unknowns -- are a challenge.

Robust Checkpointing Strategy is Key

To solve this particular type of problem, the research team designed a robust parallel solver based on a two-dimensional partitioning strategy, relying on BLAS3 computations. The size of each block in the partition depends on the size of the global problem and is computed to maximize parallel processor efficiency. The solver includes:

- an optional partial pivoting strategy
- a reordering scheme to maximize the stability of the diagonal block factorizations
- a robust checkpointing strategy

The checkpointing feature was key to the success of the run, which would have otherwise aborted. This feature is rather advanced, allowing a given processor to restart using data previously assigned to a different processor.

Details Reveal SP2's Reliability

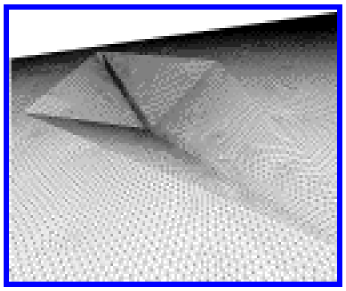

The benchmark problem used in this calculation is designed for solving radar illumination problems. The object used in the cross-section computation is a metallic winglike structure, with a total surface area 0.27 m^2 and a slot that functions as an antenna.

To compute the radar cross section at 15 GHz (or 675 square wavelengths), the solver discretized the surface into 66,716 triangular elements and 133,432 nodes, giving rise to 100,000 unknowns. A total of 5,760 right-hand sides, corresponding to 5,760 different incidences, were generated.

Using 144 of the SP2's 164 processors, the radar cross section of the structure was computed in 21 hours, 24 minutes. Specifically, the system matrix was computed in 27 minutes; the factorization took 12 hours, 27 minutes (including 65 minutes for parallel I/O operations); and forward/backward substitutions required 8 hours, 30 minutes. A total of 1.54 terabytes of data was transferred to and from the processors' local disks.

The processors sustained an impressive 29-plus gigaflops, including the interprocessor communication and I/O manipulations. Noteworthy was the reliability of the SP2, which ran for 21 consecutive hours, after some initial problems involving switch connectivity were resolved.

Behind Impressive Results -- Teamwork

Important to the success of this run was the NAS parallel systems and scientific consulting staff, without whom the simulation could not have been performed. The team worked around-the-clock: Chuck Niggley, scientific consulting group leader, coordinated the system and staff resources required. Ernst Kimler, NAS control room analyst, wrote the scripts for interrupting regularly scheduled jobs, which were required for setting up the dedicated system. Mary Hultquist, lead systems analyst for the SP2, resolved other hardware and software problems as they arose; for example, shortly after the job started, one node lost switch connectivity (a frequent problem with the SP2 during the prior week). Hultquist identified potential problem nodes and, because of the checkpointing feature, was able to move data files to different nodes and restart the job. Al Bayucan, control room analyst, worked closely with Hultquist, and also returned the SP2 to normal operations quickly once the special run was completed.

Solver Has Many Applications

Aside from radar illumination, the dense solver has many applications in industry, including: nondestructive testing, detecting cracks in a structure, and detecting/evaluating welds in printed circuit boards and computer chips. It can also be used by health care specialists for non-invasive medical visualization and diagnosis.

In a letter to NASA administrator Daniel Goldin, Irving Wladawsky-Berger, general manager of IBM's

RS6000 Division, commented: "The CAS [computational aerosciences] community will see this as a new benchmark and interpret it as further evidence that parallel processing will open new possibilities to aircraft design and testing."

Simulation in `Almost-real' Time

The Boulder simulation was an offshoot of a much larger project involving fluid/struture/thermal/electromagnetic interaction problems. The ultimate goal of the project is to someday perform a truly multidisciplinary flight in "almost-real" time. The next step is rebuilding the computation at twice its current size, which will require more nodes -- since the NAS SP2 has only 160 "wide" nodes (IBM RS6000/590s) -- although it is currently the largest wide node SP2 worldwide. The Boulder job achieved 90 percent peak performance on the wide nodes, as opposed to 60 percent on a "thin" node SP2.

For more information on the Colorado team's work, send email to **charbel.farhat@uofcolorado.edu**.

Charbel Farhat, professor of aerospace engineering at the University of Colorado, Boulder, led the research team. He holds a doctorate in computational mechanics from the University of California, Berkeley, and has received several awards, including the prestigious Aerospace Structures and Materials Award in 1994. He lectures on computational mechanics and has written over 125 publications on this topic.

Next Article | Contents | Main Menu | NAS Home

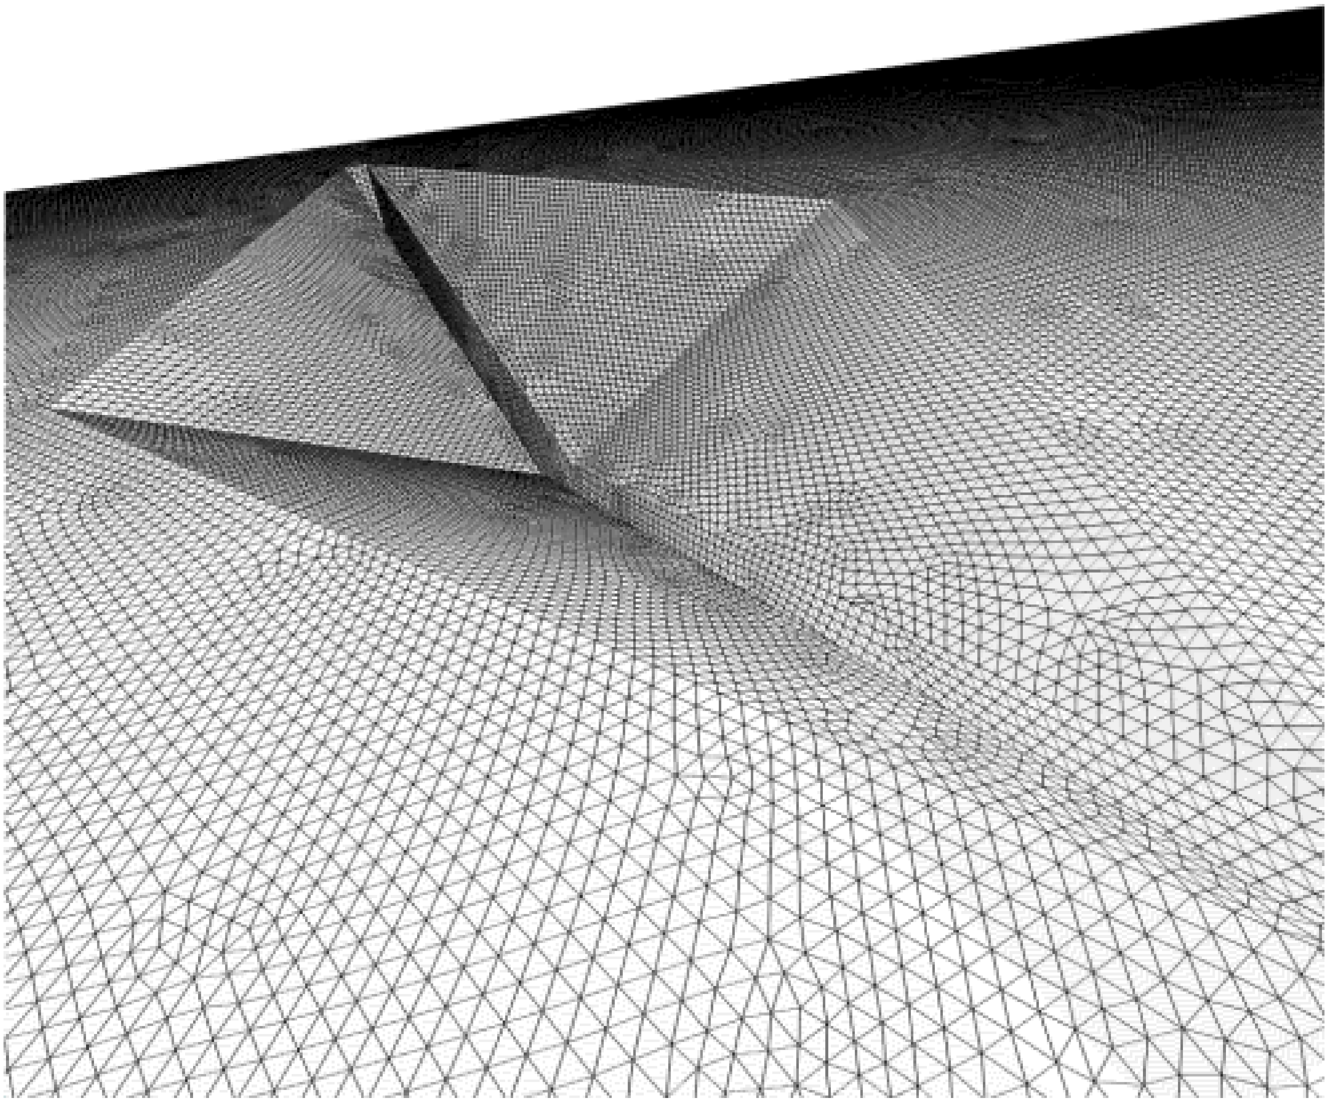

A partial view of the surface discretization of the wing-like metallic structure. The structure is part of a benchmark used in the out-of-core electromagnetic scattering computation run on the IBM SP2 at NAS. The surfacic mesh contains 66,716 triangular elements and 133,432 nodes. Rendering was performed with the TOP/DOMEC software.

to the article

The NAS Charter: 'Fundamental R&D' for NASA Information Technologies

By Elisabeth Wechsler

With the May announcement of further NASA budget cuts and personnel reductions, Dave Cooper, Director of Information Systems (which includes the NAS Systems Division) at Ames, talked with *NAS News* about the future role of the NAS Program. Cooper served as Chief at NAS until last October when he assumed his current position.

"NAS will continue to play the pathfinding role in high-performance computing environments, particularly high-speed processors (HSP), focusing on innovative long-term basic research in areas of fundamental importance to NASA missions in aeronautics and space," Cooper said. These areas include systems software, scientific visualization, mass storage, and high-speed networks. He added that the NAS Systems Division "will serve as a testbed for these new technologies, taking the development activity to demonstration projects."

Walt Brooks Named Acting NAS Chief

Walt Brooks

Cooper also announced that Walt Brooks, formerly Chief of Information Systems' Programs and Projects Division, would be "detailed to the NAS Division as Acting Chief for 120 days," pending NASA headquarters' approval of his permanent appointment. Brooks, who has worked at Ames for 17 years, previously served as Assistant Director of Aerophysics at the center. He performed research for his thesis at Brookhaven National Laboratory and received a doctorate in physics from Stevens Institute of Technology. In 1991, he was selected as a Sloan Fellow at the Stanford Graduate School of Business, where he earned a master's degree.

[Marisa Chancellor, Deputy Chief of NAS, served as Acting Chief until she took family leave in May; Chancellor plans to return in September. Eric Hibbard continues to serve as Acting Deputy Chief in her absence.]

Knowledge-based Technologies

NASA has named Ames as the Center of Excellence for Information Systems Technologies. As a result, Cooper said that his directorate will focus on "mostly knowledge-based technologies, including expert systems, artificial intelligence, neural engineering, numeric computation, simulation technology, intelligent interfaces, and software engineering.

"Initially, we're going to address those areas that meet the requirements of existing programs," he said. "Remember, however, that Ames is an aeronautics center and we have major commitments to the NASA aero program, which we will continue to support and advance."

Cooper went on to say, "We expect NAS to continue to play a pathfinding role for supercomputing within NASA. And that, coupled with our lead center role for High Performance Computing and Communications (HPCC), will allow us to continue focusing on systems software for MPP systems, which is applicable to a heterogeneous computing environment that would also include workstation clusters and conventional vector supercomputers," he added.

Systems Software More Important

Cooper strongly believes that "the success of teraflop computing does not depend on hardware right now. The key is systems software. We are the lead center for that within NASA and we're working with other government agencies, as well as with industry, to continue that effort."

He explained: "As you approach teraflop computing you get exponentially larger datasets. You've got to have software that can rapidly and effectively display meaningful results so that you can quickly analyze your data and then make the necessary changes for the next set of calculations."

With the IBM SP2 in place, NAS can now "test the new software that we're developing for new MPP systems," Cooper said. "We can't just throw these new technologies over the fence," he added. "Our customers understand and accept the use of the NAS Facility as a testbed, particularly those dealing with the MPP systems and workstation clusters. By engaging in these activities, we hope to give our customers a bit of a lead over any foreign competition that might arise."

No Big Changes for Industry Users

Cooper predicts little impact to NAS customers -- "particularly those in industry" -- from the changes taking place at Ames. He cited the continued presence of the HPCC Program, Computational Aerosciences applications, National Research and Educational Network (NREN) project, and Information Infrastructure Technology Applications (IITA).

"The NAS Facility's budget remains relatively intact. We have solid funding from the Office of Aeronautics, so we will continue to support our customers," Cooper said.

User Survey Confirms That Design Optimization Is 'Top Issue'

By Elisabeth Wechsler

An informal survey of participants at the NAS User Interface Group meeting in January confirmed that "optimizing the aircraft design process--getting the product to market sooner--is most important" to NAS aeronautics industry users, according to Bruce Blaylock, NAS Division Chief Scientist.

Survey respondents asked NAS to assist in this process by helping industry users "parallelize" their environments and "expand their computing capabilities from vector processing to Massively Parallel Processors (MPP) and clusters," Blaylock said.

More 'Parallelization'

"Even though we're doing a lot of parallelization work, we need to do more," noted Blaylock, who along with Tom Pulliam and Alex Woo (both in Ames' Advanced Computational Research Branch), designed the five-part questionnaire. "The future of computing is parallel and cluster environments. We need to help our users make that transition," Blaylock added.

Examples of NAS-based support would include system software and tools, early use of new computers, and consulting help, Blaylock continued. Another area of broader participation would be hands-on collaborative efforts, such as the DARWIN [Development Aerodynamics Re-evolutionizing Wind tunnels and Intelligent systems for NASA] project.

"High Performance Computing is not just a Cray-class computer on the floor but a way to gang all computational facilities into one system and make use of it," Blaylock said.

Services `Beyond Compute Cycles'

Two survey topics offered particularly interesting information for NAS, Blaylock commented. One of these asked whether users wanted "more services beyond compute cycles" from the NAS organization. Areas of response included parallelization and optimization of workstation clusters; system software to port code from parallel machines to shared memory multiprocessors; grid generation and output analysis; and new visualization software and services.

As part of the survey, UIG participants were also asked to complete: "NAS would help me the most if it did the following..."

Many respondents requested support for parallel code optimization and scientific visualization. A fairly representative response was: "Help remote users learn to optimize codes for operation on the HSPs [High Speed Processors], whether vector or parallel processing. Help remote users with flow visualization for transient data."

Other survey topics addressed NAS policies for scheduling large jobs, optimizing throughput, and allocating supercomputing resources, as well as queried customers about their future applications capacity requirements.

"Respondents want mostly large jobs, but not exclusively," Blaylock said.

NAStore Earns High Marks

NAStore, the NAS mass storage capability, received high marks. "I think NAS has done a tremendous job in this area in the past year," noted one survey respondent. Another commented that NAS has provided "fast, convenient access to long-term storage."

The survey, which was completed by most industry representatives attending the UIG meeting, will "give us some leads [with which] to pursue discussions over the year," Blaylock said. "We want to make the NAS environment more usable, reliable, and robust for users."

Next Article | Contents | Main Menu | NAS Home

Adapting to the Cluster Environment: Dynamic Parallel Applications

By Bill Saphir

The High Performance Computing and Communications (HPCC) Program workstation cluster located at NAS represents the latest, and hopefully, the final stage in the "de-evolution" of parallel computing -- the trend from tightly integrated parallel computers to loosely coupled cluster. In contrast to its predecessors, it is built entirely from commodity hardware and software. "Faster, cheaper, better"? Well, it's cheaper, but the jury's still out on the others.

In some ways, the HPCC cluster and its industry counterparts -- collections of workstations on people's desks -- may turn out to be fundamentally different from traditional supercomputers. The difference has to do with the "noisiness" of the environment. In a noisy system, measurements of system performance yield noisy or erratic data. A clean system produces data with little noise. Without the tight integration of parallel computers such as the Thinking Machines Corp. CM-5 and the Intel Paragon, it's difficult to control, measure, and predict every aspect of system performance.

There are many sources of noise on workstation clusters: the nodes usually run more system processes than do nodes of dedicated parallel computers; primary users may run screen savers, or may have abandoned or forgotten about other CPU-intensive jobs. Nodes running these extra processes are, in effect, slower than other nodes.

While many current parallel programs slow down or die in the face of unpredictable parallel environments, some applications can pre-empt potential problems by responding intelligently to a dynamically changing environment.

Two general categories of such intelligent response are dynamic load balancing and dynamic resource adaptation. Neither strategy is in wide use at NAS, and both can be difficult to implement, but the HPCC cluster and its industrial cousins provide a motivation and a testbed for exploration.

Static Load Balancing

Most parallel applications run at NAS are statically and evenly load balanced, that is, they split a problem into equal-sized chunks and never reevaluate this initial decision. Static load-balancing works well in a clean environment where all nodes run at a constant speed.

In workstation environments, even small fluctuations in processor performance can hurt the performance of these statically load-balanced applications. At first glance, one processor in a parallel application of size N should only affect overall performance by an amount of order 1/N. However, in a statically load-balanced application, all nodes wait on the slowest one -- so, a single slow node can slow an application by an arbitrary amount. For example, on the NAS IBM SP2, which is clean by workstation cluster standards, parallel applications running on NFS server nodes have experienced up to 50 percent slowdown.

Dynamic Load Balancing

The load balance problem can be addressed by both system administration policies and by dynamic load balancing within an application. NAS staff in the parallel systems and workstation groups are currently making every effort to avoid unnecessary system noise, for example, by rearranging the location of global file systems and prohibiting interactive logins to most of the nodes. Some noise is still inevitable, however -- especially out in the "real world," where there's typically less control over the environment -- so it makes sense to investigate dynamic load balancing.

How can you add dynamic load balancing to an application? The answer is different in every case. Applications in which a manager process hands out small independent tasks for worker processes to perform are trivial to balance dynamically. This is significantly more difficult in applications that distribute a single grid across multiple processors. To handle slowly varying or quasi-static loads, it is almost always possible to implement infrequent dynamic rebalancing. Even the most minimal scheme, such as balancing a load at start-up time (based on measured system response, not static configuration), can help reduce processing time significantly when an application competes with other long-running jobs or when it runs on an inherently heterogeneous system.

`Dynamically Adapt' to Changes

Another annoying feature of real-life workstation clusters is that they fail. A janitor may unplug a computer to get power for a vacuum, a hardware or software problem may cause a network to die. Users returning for late-night work might kill all jobs running on their workstations. On the other side of the coin, resources may become available in the middle of a calculation, when, for instance, an engineer finally decides to call it quits at 2:00 a.m. Both situations require that an application be able to dynamically adapt to changing resources. Providing fault tolerance is generally more important than making use of new resources, but the two are related.

Fault tolerance can range from checkpointing, which provides passive tolerance of failure, to failure detection and failure recovery. None of these will be provided automatically by system software in the foreseeable future.

Checkpointing is the easiest type of fault tolerance to provide. An application can checkpoint itself by saving its state to a file at regular intervals. Heavily used applications that run for more than an hour or so are prime candidates for adding checkpointing.

Active fault tolerance requires at least minimal system support. Currently, the PVM (Parallel Virtual Machine) message-passing library allows applications to detect and respond to process failure. The NAS parallel systems group generally recommends MPI over PVM for reasons of performance and standardization, but MPI does not currently support fault recovery.

The MPI Forum is currently considering dynamic process management and expects to have a draft of an expanded MPI-2 standard and trial implementations by December 1995. The parallel job management systems in use at NAS, PBS, and LSF [Load Sharing Facility], currently provide no special support for fault tolerance, but NAS is exploring the issue and would like user input. See current information about MPI.

What To Do?

The preceding discussion contains few details on adding dynamic response to parallel applications -- unfortunately, the details are different in every case. If you're interested in adding dynamic load balancing or fault tolerance to your parallel application, here's what to do:

- Contact NAS User Services with questions about where to start, and ask to speak to a parallel systems consultant. Send email to **nashelp@nas.nasa.gov**, or call (415) 604-4444 or (800) 331-USER).
- If you're working with MPI and need specific functionality, contact a NAS representative on the MPI Forum, who will advocate for you (send email to **wcs@nas.nasa.gov**).
- Let us know about any system software support that would make it easier to support dynamic response by contacting NAS User Services.

Next Article | Contents | Main Menu | NAS Home

Parallel Tools Meeting Joins Users, Developers, Researchers

By Marcia Redmond

Approximately 100 parallel tool users and developers from organizations that included IBM, Sandia National Laboratories, and the Institut fur Informatik, Munich, gathered at the second annual Parallel Tools Consortium (Ptools) Meeting. Sponsored by the NAS Systems Division, the meeting was held at NASA Ames Research Center, May 22-24.

Participants were welcomed by Jerry Yan, NAS researcher and Ptools Consortium Steering Committee member, and Eric Hibbard, NAS Acting Deputy Division Chief. "We believe that this type of forum is very important -- not only to the NAS Systems Division, but to Ames and NASA in general -- for technology transfer and collaborative work," Hibbard said. The event included small group sessions, key speaker presentations, a panel session, and hands-on demonstrations.

Cherri Pancake, a professor at Oregon State University (OSU) and Consortium Chair, said that "the Consortium is a forum for interaction between three groups -- tool users, tool implementors, and tool researchers." Key presentations focused on the results of collaborative projects by universities, industry, and government labs.

One of these projects -- a joint effort by Ames, OSU, and Intel -- is the Message Queue Manager (MQM). According to Don Breazeal, user environment group manager at Intel, MQM "provides an intuitive way to examine the message-passing state."

One component of MQM, the Standard Query Interface (developed at OSU), is located between the graphical front end, which manages the filters, and the target system. "It is the responsibility of the implementor of the target system to implement this query interface and provide the information that MQM is asking for," Breazeal said. He noted that Intel has successfully replaced the Paragon's debugger with MQM on the front end.

Ames plans to modify the PVM-AIMS [Parallel Virtual Machine-Automated Instrumentation Monitoring System] software, developed at NAS, to support the Standard Query Interface.

Rusty Lusk, senior computer scientist at Argonne National Laboratory (ANL), presented "Scalable Unix Commands on Parallel Computers," a collaboration with Jim Cownie (Meiko Inc., UK), Dan Frye (IBM's Power Parallel Systems Division, Kingston, NY), Lusk, and Bill Gropp (also of ANL). The team developed a set of commands because, as Lusk explained, "the users need simple, familiar ways to copy a file to the local file source on each node, find all [their] processes running on all nodes, and test a condition on all nodes without getting swamped with output." One example is a command, named "pfps," which sends mail to everyone (other than root) on the system that has a long-running job, including the name of the job, the amount of time, and the processor where the job is running.

For more information on the Ptools Consortium, send email to **yan@nas.nasa.gov**. Details on current Ptools projects are available online.

Next Article | Contents | Main Menu | NAS Home

 of Tools for CRAY C90 Users *By* R.K. Owen

The NAS and Aeronautics Consolidated Supercomputing Facility (ACSF) environment provides several math libraries to help users develop code on the CRAY C90s. These libraries provide a large collection of mathematical, statistical, and special function subprograms that can be readily used from within a Fortran or C program. The libraries have been well tested and provide subprograms for common numerical operations, which saves much development time and source code size.

Typically, these libraries offer subprograms for: solving linear systems, eigensystems, and interpolation; evaluating quadrature, differential equations, and transforms; and performing special function, nonlinear optimization, regression, correlation, and random number generation. Some have specialized routines for solving linear systems or Fast Fourier Transforms (FFTs) out-of-core (that is, not entirely in memory), sparse linear solvers, or Helmholtz partial differential equation solvers.

Online information or interactive document facilities are available for some of the libraries. Look at the *libraries* (*local*) man page for what's available on the CRAY C90s, vonneumann (NAS) or eagle (ACSF). Use *alldoc*, a keyword search tool available on both machines, to find subroutines and functions. The libraries listed below are available on both machines, unless noted. For more detailed information, see the NAS Scientific Consulting (SciCon) math library page.

IMSL 2.0 is `Painless'

Visual Numerics (formerly International Math and Statistics Library) provides general math, statistical, and special function subprograms -- and does so in as painless a way as possible. It provides the easiest practicable interface for users, handles most storage needs internally, and has a special interface if internal storage is insufficient. Routine names (constrained to eight characters) often indicate its purpose. Some of its unique capabilities are: a full range of B-spline interpolation and approximation routines, cubic splines with periodic boundary conditions, a full range of special functions related to the gamma function with real and complex arguments, and the inverse error function, various routines for convolution and correlation, some special integrations such as Fourier integrals. It also includes a slate of unique statistical routines for doing discriminant analysis, cluster analysis, density and hazard estimation, exploratory data analysis, empirical probability distribution analysis, and has several data display routines.

An easy-to-use online interactive document facility called *imsl.doc* is available on eagle.

NAG is `Expansive and Dynamic'

Numerical Algorithms Group (NAG) provides many of the same general and special function subprograms as IMSL, but is more expansive and dynamic. The mix of routines constantly changes from one version to the next; deprecated routines are tagged for future removal and obsolete ones are deleted. However, the functionality is never lost -- only the interface is changed. Routine names are somewhat cryptic, but are consistent within the NAG framework. Some unique capabilities are: more special functions for complex variables, integral equation solvers, linear algebra operations for non-square matrices, orthogonalization, robust regression, and a full suite of generalized linear models to choose from.

The NAG library was recently upgraded on the CRAY C90s and has been made more efficient by removing the NAG LAPACK and BLAS equivalents that coincide with routines available in the Cray Research Inc. native scientific library, *libsci*. (A typical example is the matrix-matrix multiplication routine F06YAE, which is identical to the BLAS routine SGEMM.) Using the *libsci* BLAS routine results in a two to five times speed-up over the NAG routine, depending on the matrices sizes. The library itself was modified so that even internal calls now use the more optimized Cray libraries when available.

NAG has an online document browser called *nag.doc* on eagle. It's not as easy to use, but is up to the task of finding a specific routine that resolves a particular problem. See the man page *nag.doc* for useful tips.

SLATEC 4.1 is `Publicly Available'

Sandia-Los Alamos-Air Force Technical Library is a collection of public domain general math, statistical, and special function subprograms that have been packaged, documented, and tested. The source code is available from **netlib@ornl.gov** (just send the message "help" for more information). The sources are also available locally in /usr/local/src/lib/slatec4_1.

Some public domain libraries included are: BLAS, EISPACK, FISHPACK, FNLIB, LINPACK, PCHIP, QUADPACK, SLAP, and FFTPACK. Like the NAG library, however, the BLAS, EISPACK, and LINPACK routines that duplicate those in libsci have not been included in the SLATEC object library.

SLATEC has an online document facility, *slatec.doc*, which is available on both vonneumann and eagle, that searches for routines by keywords or by index. It has been modified locally to be easier to use and browse.

Use BCSEXT 3 for Large Problems

Boeing Computer Service provides real and complex linear systems and eigensystems using out-of-core methods, real and complex sparse linear systems and real sparse generalized eigensystems, and two-dimensional (2D) and three dimensional (3D) complex FFTs using out-of-core methods. An older general library named BCSLIB is no longer supported by Boeing and should no longer be used for code development -- there is no guarantee that it will be available on future platforms.

The Cray *libsci* includes a virtual LAPACK (see the man page *vlapack* for more details) that provides a real out-of-core linear solver and a real or complex virtual BLAS for matrix-matrix operations. However, the BCSEXT library is easier to use and covers more types. Online documentation consists of PostScript files in /usr/local/src/lib/bcsext/doc/. The directory also contains plain, loosely formatted text files that are useful when searching for keywords.

Crayfishpak for Helmholtz Equations

Green Mountain Software provides a successor to the public domain FISHPAK library of Helmholtz partial differential equation solvers.

Crayfishpak uses finite difference and FFT methods for solving 2D and 3D geometries for cartesian, polar, cylindrical, surface, and spherical coordinates. This library is only available on vonneumann, and online documentation is located in /usr/local/src/lib/crayfishpak/doc/. The file "Paper.ps" is a PostScript file that gives a general overview; the rest are text files that cover each of the driver subprograms. Example code for each are found in /usr/local/src/lib/crayfishpak/example/.

Library Comparisons

The alldoc command on the CRAY C90s is useful for finding which routine in any of the libraries can solve a given problem. Use the **-K *keyword*** option to start a keyword search, then use **-F *keyword*** for subsequent searches that restrict the scope to subprograms that have satisfied the previous keyword searches. An example of this gives the search results shown in the table below, which finally zeros in on two routines: NAG's f06sde or SLATEC's chbmv. Type the subprogram name on the command line after alldoc (for example, **alldoc f06sde**) for detailed information on the given subprogram.

alldoc command	entries found
-K complex	581
-F matrix	288
-F hermitian	102
-F band	16
-F vector	2

Suppose all of these libraries offer routines to solve the given problem; for example, you need to find a vectorized uniformly distributed pseudo-random number generator. See Math library CPU timings for selected routines across a broad range of applications. Looking through this web page, you can find a section that covers this topic with raw CPU times listed and the option to examine either PostScript or GIF files in color or black and white. This gives the maximum latitude to view or print the plot files.

Here's an example of using these libraries from within a Fortran or C program.

If you have questions or need help with the math libraries on either vonneumann or eagle, call NAS User Services at (415) 604-4444 or (800) 331 USER.

Next Article | Contents | Main Menu | NAS Home

NAS Researchers Win Allen Award

David H. Bailey, and four other NAS-based scientists, will receive the 1995 H. Julian Allen award at an upcoming awards ceremony for two research papers about the NAS Parallel Benchmarks, published in 1991 and 1993, respectively. The award, given annually by NASA Ames Research Center for work done primarily at the center, is accompanied by a $5,000 stipend, which will be split among the 13 authors.

Besides Bailey, those authors at NAS include Eric Barszcz, Rod Fatoohi, Tom Lasinski, and Sisira Weeratunga; Robert Schreiber is at the Research Institute for Advanced Computer Science. Other authors no longer at NAS or Ames include John Barton; David Browning (Technische University, Netherlands); Russell Carter (Geli Engineering); Leo Dagum (Silicon Graphics Inc.); Paul Frederickson (Math Cube Associates); Horst Simon (Silicon Graphics Inc.); and V. Venkatakrishnan (ICASE).

The specific submission cited for this year's award consisted of the NASA Technical Report, "The NAS Parallel Benchmarks," which was first released in 1991. A condensed version of this paper was published in the *International Journal of Supercomputer Applications*, vol. 5, no. 3 (Fall 1991), p. 66- 73. "NAS Parallel Benchmarks Results" was published in *IEEE Parallel and Distributed Technology*, Feb. 1993 (premier issue), p. 43-51.

Testbeds Explore Transition From AEROnet to ATM

By Elisabeth Wechsler

Two testbeds involving AT&T and U.S. Sprint, as well as six NASA High Performance Computing and Communications (HPCC) Program centers, were set up by NAS this spring to test Asynchronous Transfer Mode (ATM) network service. The objective, according to John Lekashman, manager of the NAS network group, is to eventually migrate the AEROnet and NASA Science Internet dedicated backbone infrastructure, as well as other NASA networks, to a shared ATM service infrastructure.

"This involves providing an access connection to the ATM service rather than the more costly cross-country dedicated circuits, which are used today," explained Anthony Lisotta, a member of the project team.

Better Throughput Plus Cost Savings

In addition to cost savings of approximately 30-40 percent, ATM service is expected to greatly increase network throughput capacity between NASA HPCC centers, Lekashman pointed out. The current AEROnet delivery bandwidth is 45 megabits per second (mb/sec). This throughput rate is expected to increase threefold to 155 mb/sec after the first phase of testing is completed this fall.

One testbed links NASA Research Centers (Ames, Langley, and Lewis) and Marshall Space Flight Center, using AT&T's data system. The other testbed connects Ames, Langley, Lewis, Jet Propulsion Laboratory, and Goddard Space Flight Center via U.S. Sprint's data system, which makes up the NASA portion of the National Research and Educational Network (NREN). Plans for an additional testbed (involving MCI, Ames, Boeing, and McDonnell Douglas) are being discussed, Lekashman said.

NAS is working with Marshall's Program Support Communications Network office on the AT&T testbed. The Department of Energy (DOE) office at Lawrence Livermore National Laboratory has contributed "knowledge based on past experience" to the U.S. Sprint testbed. The DOE has simultaneously deployed the Energy Sciences Network over the same ATM service, Lekashman said. "That system is now in full service at 45 megabits per second for DOE."

Data Collection Began in May

Since May, the NAS team has been collecting data on throughput, error, and delay, as well as security paradigms. Development versions of applications such as Distributed Virtual Wind Tunnel (dVWT) were scheduled for testing last month, Lisotta said. "We also want to run real-time video for the DARWIN project, of which one component is the Remote Access Wind Tunnel project," he added. (DARWIN is the "Development Aerodynamics Re-evolutionizing Wind tunnels and Intelligent systems for NASA.")

The phasing of AEROnet to ATM is expected to begin by year end. "We're trying to make the transition as seamless as possible," Lisotta noted.

Whether users notice any impact on their network "will depend on the specific site," Lekashman explained. "If the site's local area network infrastructure is only ethernet, they won't see any performance increase. Those with HiPPI, FDDI, and ATM `to the desktop' will."

The NAS team (which also includes Jude George, Arshad Khan, Alfred Nothaft, and Kelly Russell) will examine "interoperability issues" between different long-distance communications carriers, Lekashman said. "Maybe one system will be chosen based on cost--or perhaps different systems in different parts of the country will end up costing less, if they are easily interoperable."

For more information, send email to **lekash@nas.nasa.gov** or **lisotta@nas.nasa.gov**.

CRAY C90 Upgrade Gives ACSF Users More System Time, Better Turnaround

Recent improvements to the Aeronautics Consolidated Supercomputing Facility (ACSF) CRAY C90 supercomputer have given users more space to run jobs and have eliminated system idle time during business hours. Changes include increased memory, a new multitasking queue, and larger queue memory limits for all jobs. The upgrade was installed in April.

Responding to user requests, the NAS high speed processor (HSP) support group upgraded the C90 from 128 megawords (MW) of memory to 256 MW, and increased job size limits for all queues from 32 MW to 64 MW, resulting in faster job turnaround.

"The NAS Systems Division was able to reconfigure the ACSF hardware to significantly reduce cost while maintaining the same capabilities," said Eric Hibbard, NAS Acting Deputy Division Chief. As a result, Hibbard explained, "we were able to buy the memory upgrade and stay within budget. This change should make the system much more appealing to ACSF customers who want to run large memory jobs."

Unlimited Submission of Multitask Jobs

The new multitask queue (named mtask) runs 24-hours a day and gives users up to 10 hours of CPU time for each job; single-task jobs have a limit of 8 hours. The memory limit for another queue (mtasklg) reserved for large multitask jobs has increased from 96 MW to 200 MW and allows jobs to run up to 16 hours. Either of these queues can, if necessary, allocate all eight CPUs of the C90 to a single job. There is no limit on the number of jobs users can submit to these queues, and users can control the order of their jobs in the queue by using the *qorder* command.

`Submit More Jobs'

Alan Powers, who leads the HSP support group, encouraged ACSF users to "submit more jobs--at any time--and more multitasking jobs." Even though each user can submit an unlimited number of jobs, these are selected from an eligible pool, and each user is allowed up to three jobs at a total of 96 MW of memory and 14 hours of CPU time. As jobs run, other jobs automatically migrate from ineligible to eligible states. Users should contact their local resource monitors to request additional computer time.

About 600 users run computations at the ACSF, which was created in 1994 as part of the effort to decrease government spending. Previously, each of three sites--NASA's Ames, Langley, and Lewis Research Centers--operated and maintained its own supercomputing facility. Each site now pays for a portion of the consolidated facility, with an estimated savings of $30 million over five years.

Next Article Contents Main Menu NAS Home

Home of the Future

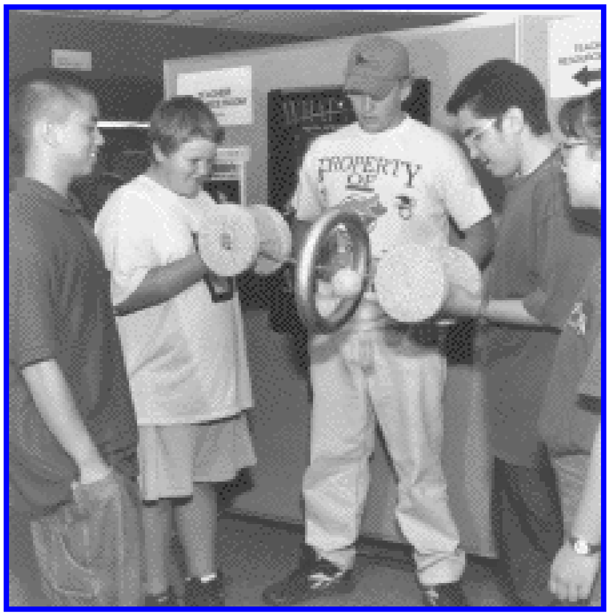

Model space colony designed by ninth-grade students at John Swett High School, Crockett, CA, who shared the grand prize in the 1995 Space Settlement Contest with students from Live Oak High School, Morgan Hill, CA. This year, 145 students participated--more than twice the number of entries last year. Students from grades 6-12 were required to submit reports and supporting material. Entries were judged on technical and artistic merit. All contestants received a letter of commendation, and students living nearby toured NASA Ames Research Center on May 23.

The idea behind the contest, begun three years ago, was to "promote NASA's science education effort" and to "train the next generation of engineers to build space settlements," said scientist Al Globus, who conceived the contest and co-judged entries with colleague Kristina Miceli (both are in NAS's Data Analysis Branch). Lesson plans designed by Tug Sezen, a teacher at John Swett, will be available through the Ames Teacher Resource Center before the start of the 1995-96 school year.

- Next year's contest or, send email to **globus@nas.nasa.gov**.

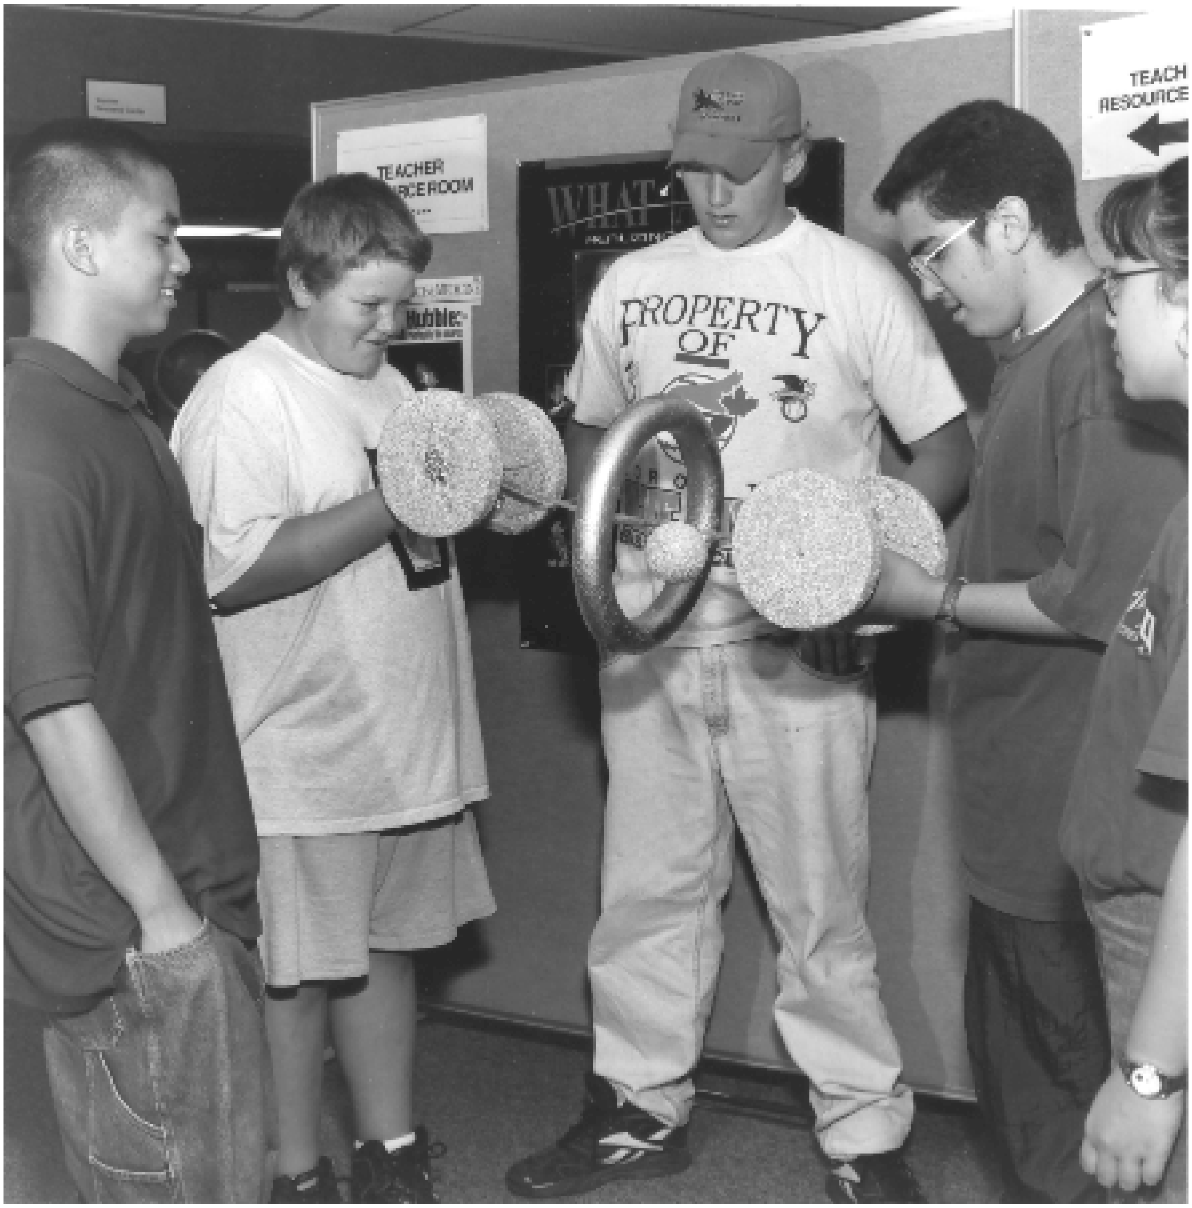
TEACH
RESOURCE
PROPERTY
OF

Next Article | Contents | Main Menu | NAS Home

July - Aug 1995, Vol. 2, No. 11

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: Charbel Farhat, Mike Gerald-Yamasaki, Kristina D. Miceli, R. K. Owen, Marcia Redmond, Bill Saphir

Image Enhancements: Chris Gong

Other Contributors: Chris Atwood, David Bailey, Tim Barth, Bob Becker, Bruce Blaylock, Don Breazeal, Walt Brooks, Nick Cardo, Dave Cooper, Jim Craw, Nicolas Durante, Pat Elson, Dave Faust, Robert Garza, Al Globus, Bob Haimes, Eric Hibbard, Philippe Homsi, Ed Hook, Mary Hultquist, David Kenwright, David Lane, John Lekashman, Anthony Lisotta, Terry Nelson, Alan Powers, Leigh Ann Tanner, Dani Thompson, Sam Uselton, Tom Whittaker, Alex Woo, Jerry Yan

NAS NEWS

Volume 2, Number 11 July – August 1995

INSIDE: A Special Section on Unsteady Visualization – see page 5.

[illegible]

Large, Out-of-Core Calculation Runs on the IBM SP2

by Charbel Farhat

[illegible]

[illegible]

Robust Checkpointing Strategy is Key

[illegible]

- [illegible]
- [illegible]
- [illegible]

[illegible]

Details Reveal SP2's Reliability

[illegible]

Continued on page 3

This Issue:

The NAS Charter: 'Fundamental R&D' for NASA Information Technologies

by Elisabeth Wechsler

[illegible]

Walt Brooks Named Acting NAS Chief

[illegible]

Knowledge-based Technologies

[illegible]

Systems Software More Important

[illegible]

Continued on page 11